

I'm Tom Armour and I am delighted to be here in beautiful Dallas in what still seems like August to tell you about the new Asymmetric Threat Initiative of DARPA's Information Systems Office.

I can still remember a metaphor Jim Woolsey used in a statement he made after being named to be the next Director of Central Intelligence.

This was a few years after the Berlin Wall had come down and the Soviet Union had disintegrated.

People then were talking giddily about a "peace dividend."

Woolsey pointedly said that while the "big, bad bear" was gone, the woods were still filled with lots of "poisonous snakes" and other little nasties and was every bit as dangerous a place as it had been ... if not even more perilous.

Woolsey was referring, in part, to what we now call the asymmetric threat.

Let me briefly talk about how these threats are different from Woolsey's big, bad bear.

Conventional military threats --an armored division, a naval task force, or an air wing, for example-- are easily observed by reconnaissance satellites and other technical collection assets.

If the unit should take hostile action, attribution is easy.

Entry costs are high: the necessary capabilities take a long time to build and are very expensive.

The intelligence analyst whose job it is to monitor these threats can do so by consulting a few key collection sources.

And she can predict the potential range of actions using models of the physical kinetics and kinematics that govern the behavior of conventional threats.

But the new asymmetric threat is physically small -- perhaps even just a single person -- and not easily observed, especially by our existing reconnaissance systems.

Entry costs are low.

For instance, it no longer requires the resources of a nation state to develop weapons of mass destruction.

The intelligence analyst will need to consult vast amounts of information, from both classified and open sources, to piece together enough evidence to understand their activities.

And to predict the potential range of actions, the analyst will need to model the group's beliefs and behavior patterns.

The Asymmetric Threat Initiative consists of six projects within ISO:

* Human Identification at a Distance is developing new image understanding techniques to uniquely identify humans at distances of 10 to 150 meters.

* Evidence Extraction and Link Discovery is developing new information extraction and data mining technology to automatically discover and relate evidence of threatening activity from vast amounts of data.

* Wargaming the Asymmetric Environment is developing new modeling techniques to enable predictive modeling of asymmetric groups and their behavior.

* Project Genoa is developing new collaborative reasoning and structured argumentation techniques to improve threat understanding and decision-making.

* Rapid Knowledge Formation is developing new knowledge acquisition technology to enable technical experts to create comprehensive knowledge bases on any topic.

* Agent-Based Computing is developing a new technology to enable the development and use of large collections of collaborating software agents.

I will now talk in more detail about the newest of these projects, the first three on the list.

The goal of the HumanID program is to develop a system that can identify humans as unique individuals (although not necessarily by name) at a distance, at any time day or night, during all weather conditions, even with non-cooperative subjects, possibly disguised and amidst a group. This capability will be an enabler for early warning of and protection against some asymmetric threats.

Current human recognition methods require that biometric signatures be acquired from cooperative subjects in contact with or in close proximity to the sensor.

Moreover, current systems use a single biometric signature for recognition.

The HumanID program approach is to extend biometrics technology along numerous fronts.

Much of the program's technology development is focused at advancing the state of the art in facial, gait, and iris recognition through unique new sensors, unique ways of using ordinary sensors, and advanced algorithms.

The program also has research in other areas including fusion of multi-modal biometric signatures, investigating multi- and hyperspectral facial signatures, gait analysis, radar signatures, developing 3D face and body models for use with detection systems, and experimental studies of the psycho-physics of the human visual system.

The next program in the Asymmetric Threat Initiative is EELD.

The basic idea of EELD is to build a system that learns to automatically extract and then correlate patterns of evidence of threatening activities from large volumes of text.

The system would first search through large text collections, such as email messages and web pages.

It would then apply information extraction techniques to extract computer readable descriptions of facts found in those documents and correlate those extracted facts into linked patterns.

The system would then search through the linked patterns to discover novel patterns for presentation to the analyst for his feedback.

The analyst would highlight patterns of interest, correct the system's mistakes, and identify new patterns for the system to recognize.

Based on this feedback, the system would use a variety of machine-learning techniques to adjust its search, extraction, and link-discovery parameters to improve its ability to extract patterns of interest.

In the next few slides, let's walk through an example as it might look to the analyst.

After searching through large volumes of textual data, the system would select a portion of text, such as the sentences shown in the upper left.

(ANIMATION) It would then apply information extraction techniques to extract a computer-readable description of the relational facts found in that sentence and automatically store those relational facts in a database.

Those facts could then feed a link analysis tool like the one in the bottom right of the slide for further analysis by the human or the machine.

Current information extraction technology, as tested in the DARPA-sponsored Message Understanding Conferences, is able to extract simple relations from text with an accuracy of approximately 70%.

One of the technical goals of this project is to raise the extraction accuracy, for simple relations, to 90%.

The next step is for the system to discover relational-- or linked -- patterns of interest among the extracted facts.

This slide shows a small sample of the relational facts which might have been extracted prior to the bombing of the Nairobi Embassy in Africa from intelligence messages, reports, and other information sources.

(ANIMATION) The system would search through millions of such relational facts and automatically classify patterns of interest. The analyst might previously have shown the system examples of terrorist organizations and pre-operational staging.

Based on these examples, the system would have learned a new classification model for recognizing similar patterns in the future.

As the system analyzes newly extracted facts, it would recognize patterns that it would classify as, in this example, a terrorist organization or as pre-operation staging.

The system would then present these newly classified patterns to the analyst for verification.

Current machine learning and data mining technology can learn to classify objects based on the attributes of an individual object, but not based on relational patterns among objects.

Thus, another technical goal of this project is to develop new machine learning techniques which can recognize these types of relational patterns from a very few training examples.

The final component of EELD is a user-interface that would enable the analyst to correct and guide the system by annotating the results of the link discovery process. (ANIMATION)

The analyst might annotate the diagram with feedback to find more information about a particular pattern he finds interesting. (ANIMATION)

The analyst could also tell the system when it has made a mistake.

Here he is informing the system that a particular fact is incorrect because it was extracted from a bad information source - which should not be trusted in the future.

(ANIMATION) And most importantly, the analyst also could teach the system to recognize new patterns.

The last program I'll tell you about today is called Wargaming the Asymmetric Environment.

WAE is developing new modeling techniques to enable predictive and emulative modeling of the behavior and decision-making of groups posing an asymmetric threat.

Over the last ten years the Internet has enabled a virtual explosion in the availability of information, including information about asymmetric adversaries' political and military goals, organizational structure, leadership, and past attacks.

Even given this information, however, WAE does not suppose that one can accurately predict that a specific group will attack a specific place at a specific date in a specific manner. But WAE does hypothesize that we now can develop predictive and emulative models tuned to specific asymmetric adversaries.

(ANIMATION) Therefore, WAE's objective is to support a decision-maker's ability to rapidly understand the decision space by improving the

predictive focus of indications and warnings on asymmetric threats, and by emulating the full range of their behaviors.

To meet this objective we require predictive models that can reflect a representative range of an adversary's responses and that are sensitive to the behaviors and events that trigger particular alternative responses.

Today's models typically are rule, or optimization-based, and lack the agility to accurately reflect the less structured, more fluid and complex decision-making and planning of our asymmetric adversaries.

WAE hypothesizes that we can identify and model the behavioral range and triggers with a sufficient level of predictive accuracy to be used by the Operational Community.

WAE's technical approach will include empirically deriving the factors underlying the behavior and decision-making of asymmetric groups.

These include behavioral factors, intrinsic ones -- personality, leadership style, and cognitive style for example -- and extrinsic environmental influences such as political, cultural, and economic factors.

In addition we will evaluate a variety of predictive technologies to support modeling asymmetric threat behavior and decision-making.

The program also will assess the concurrent and predictive validity of the models we build.

In addition to this predictive modeling, operational wargaming entails the added requirement to generate the combined plausible ranges of interaction of adversaries, neutrals, and shifting alliances that characterize the complex and dynamic asymmetric environment.

Today's approaches have not demonstrated the predictive validity required.

WAE will develop technologies for emulating the behavior of multiple entities, with different goals, and their respective interactions in a single wargaming environment. And WAE will empirically assess the concurrent and predictive validity of this environment.

Both EELD and WAE will issue BAA announcements this fall.

Well, thank you very much for listening to me today.

If any of this has engaged your imagination, as I hope it has, I encourage you to get in touch with me or other ISO program managers to discuss your ideas.

Getting the benefit of your knowledge, expertise, and creativity is a major goal of DARPA Tech and I am looking forward to interacting with you during this conference.